Direct Differentiation Method: DDM

In the context of sensitivity analysis, most of the documents posted on this website deals with the direct differentiation method. The objective of this method is to calculate the derivative of a response, denoted by u, with respect to some parameter, denoted by θ. Structural analysis is the context, so the response may be the time-varying response, $u(t)$, at a selected degree of freedom from a linear or nonlinear dynamic finite element analysis. Although θ is one single parameter in these sensitivity-related documents, we often seek the gradient vector $\partial u/\partial \mathbf{x}$, where each entry in that vector is the response sensitivity $\partial u/\partial \theta$. In other words, the symbol θ employed here to explain the direct differentiation method is one of the parameters collected in the vector $\mathbf{x}$ appearing in other documents, such as those for reliability analysis. The following sections outline some characteristics of the direct differentiation method.

One-time Cost

Although the direct differentiation method is as efficient as they come, there is some computational cost associated with the calculation of response sensitivities. However, the most "intimidating" cost must be paid even before the calculations start. That is the one-time cost of differentiating the response equations that are implemented in the finite element code and implementing those derivative equations alongside the code for the response. The documents posted on this sensitivity webpage are intended to provide assistance. Once that work is done, and the code is debugged, all subsequent analyses produce response sensitivities in an efficient manner.

Exact

When you differentiate x^3 to get $3x^2$ then that derivative is exact. That simple observation carries over to the direct differentiation method. Because we analytically differentiate the equations of the algorithm that calculates the response, we get exact results for $\partial u/\partial \theta$. In other words, the direct differentiation method gives algorithmically consistent exact response sensitivities.

Efficient

The direct differentiation method is far more efficient than the finite difference approach, which is the primary competitor. The basic version of the finite difference method illustrates that point; consider the estimate of the response sensitivity written

$$\frac{\partial u}{\partial \theta} \approx \frac{u(\theta + \Delta\theta) - u(\theta)}{\Delta\theta} \tag{1}$$

where $\Delta\theta$ is a selected perturbation relative to the parameter value θ. The finite difference estimate in Eq. (1) has two downsides. First, it requires two separate finite element analyses in order to obtain $u(\theta+\Delta\theta)$ and $u(\theta)$. Second, it is unclear what value for $\Delta\theta$ to use. A perturbation too small may cause problems related to the numerical precision. A perturbation too large could also give an inaccurate estimate of the local derivative $\partial u/\partial\theta$. The problem is illustrated in Figure 1. The error is exhibiting a "bathtub" behaviour because both a too small or a too large $\Delta\theta$ may cause error in the estimate of $\partial u/\partial\theta$.

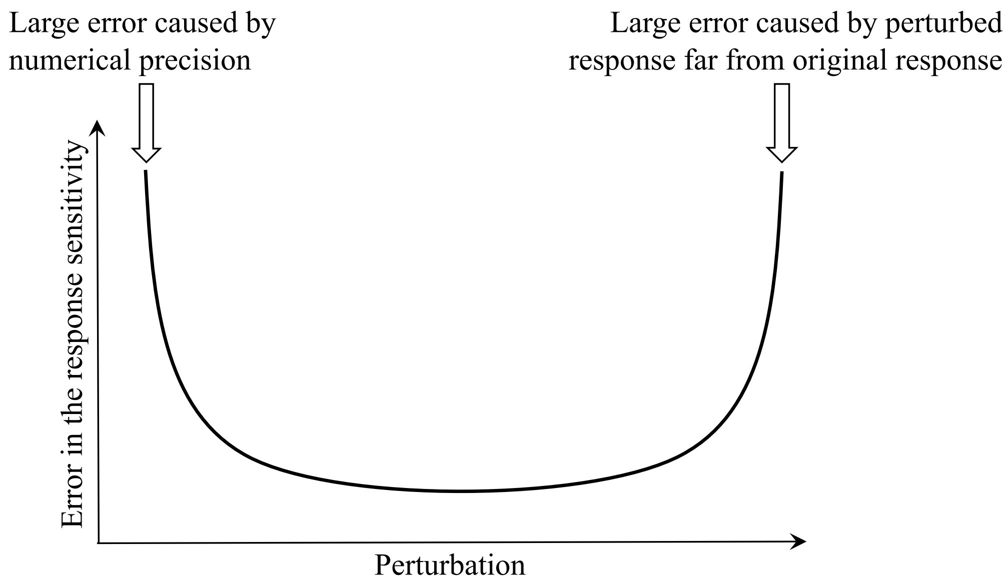

Figure 1: The "bathtub" behaviour for error in the finite difference method.

Works for All Analysis Types

Regardless of which algorithm is employed to calculate the response, that algorithm can be differentiated. It may take some time and effort, but it can be differentiated. That means the direct differentiation method works for all structural analysis types:

- Linear static analysis
- Nonlinear static analysis
- Linear dynamic analysis
- Inelastic dynamic analysis

Linear System

A key advantage of the direct differentiation method is that the response sensitivity, $\partial u/\partial\theta$, is obtained by solving a linear system of equations. The most basic linear system in computational structural analysis is the equilibrium equations $\mathbf{Ku}=\mathbf{F}$, where $\mathbf{K}$ is the stiffness matrix, $\mathbf{u}$ is the vector of degrees of freedom, and $\mathbf{F}$ is the load vector. A variety of algorithms exist for solving such linear systems of equations. In the majority of cases, the response sensitivity vector, $\partial\mathbf{u}/\partial\theta$, is obtained by solving a linear system of equations with exactly the same coefficient matrix. (The system is always linear, but for inelastic dynamic analysis with certain damping models the coefficient matrix must be modified.)

Adjoint Method for Linear Problems

This method, explained in another document, may clarify a confusion prompted above with the mixing of $\partial\mathbf{u}/\partial\theta$ and $\partial u/\partial\theta$. Oftentimes, we are looking for the sensitivity of the displacement along a specific degree of freedom, i.e., $\partial u/\partial\theta$. The adjoint method addresses that situation, by solving a single system of equations once and for all, in order to subsequently calculate $\partial u/\partial\theta$ for an array of parameters θ, simply by executing dot products. That circumvents the calculation of the full vector $\partial\mathbf{u}/\partial\theta$. However, whenever inelastic materials enter the analysis, and they usually do in nonlinear static and dynamic analyses, then the full vector $\partial\mathbf{u}/\partial\theta$ is needed at every increment. That is because the derivative of "history variables" in the materials must be calculated and stored, and those calculations require the full vector $\partial\mathbf{u}/\partial\theta$.

Algorithmically Consistent Tangent

It is mentioned earlier that the coefficient matrix in the system of equations for response sensitivities is usually identical to the coefficient matrix in the system of equations for the response itself. On this note, caution is required in nonlinear analysis, static or dynamic. In that case, there are several options for the stiffness matrix used by the Newton-Raphson algorithm that iterates to equilibrium. One option, referred to as Modified Newton-Raphson, is to use the initial stiffness matrix at each increment in those iterations. That does not work for the subsequent response sensitivity calculations. In order to obtain correct response sensitivity results with the direct differentiation method it is paramount that the tangent stiffness at the converged state is employed to solve for response sensitivities. In addition, it is important that the algorithmically consistent tangent is used in that system of equations. This is a reason why the Bouc-Wen material model is included on this sensitivity-related webpage; the difference between the continuum tangent and the algorithmically consistent tangent manifests in faster/slower convergence of the Newton-Raphson algorithm but also in exact/inexact response sensitivity results.

Explicit and Implicit Dependence

Consider a nonlinear analysis, either static or dynamic. On this website, the internal resisting forces of the structure are kept in the vector $\tilde{\mathbf{F}}$. Because $\tilde{\mathbf{F}}$ is implicitly dependent on θ via $\mathbf{u}$, and potentially explicitly dependent upon θ via the algorithm that evaluates $\tilde{\mathbf{F}}$, the derivative of the internal resisting forces reads, via the chain rule of differentiation:

$$\frac{\partial\tilde{\mathbf{F}}}{\partial\theta}=\frac{\partial\tilde{\mathbf{F}}}{\partial\mathbf{u}}\cdot\frac{\partial\mathbf{u}}{\partial\theta}+\left.\frac{\partial\tilde{\mathbf{F}}}{\partial\theta}\right|_{\mathbf{u}\text{ fixed}} \tag{2}$$

where $\partial\tilde{\mathbf{F}}/\partial\mathbf{u}$ is the aforementioned tangent stiffness matrix. The derivative $\partial\tilde{\mathbf{F}}/\partial\theta$ calculated for fixed displacements is readily obtained by differentiating the algorithm that calculates $\tilde{\mathbf{F}}$. However, two comments are attached to the calculation of $\left.\partial\tilde{\mathbf{F}}/\partial\theta\right|_{\mathbf{u}\text{ fixed}}$, each addressed in their own section below.

Fixed u not ε

The first comment related to the previous section is that $\partial\tilde{\mathbf{F}}/\partial\theta\big|_{\mathbf{u}\text{ fixed}}$ is not necessarily evaluated for fixed section deformations and strains, although the displacements are fixed. Consider the hierarchy of a typical state determination, explained elsewhere on this website, where an element takes trial displacements from the Newton-Raphson algorithm and passes them to its cross-sections. In turn, each fibre-discretized cross-section gives trial strains to each of its materials. The point in this section is that the derivative $\partial\tilde{\mathbf{F}}/\partial\theta\big|_{\mathbf{u}\text{ fixed}}$ does not necessarily translate into derivatives of section deformations being zero or to $d\varepsilon/d\theta$ being zero at the material level. As an example, consider the contribution from material fibre number j to the section forces in a frame element:

$$\tilde{\mathbf{F}}_{s,j} = \mathbf{T}_{ms,j}^{\mathrm{T}} \cdot A_j \cdot \sigma_j \tag{3}$$

where $\mathbf{T}_{\mathrm{ms}}$ is the transformation vector from the material to the section, A is the area of the fibre, and σ is the stress in the fibre. The conditional derivative of that component of the internal force is, according to the previous section, albeit incorrectly, here using the product rule of differentiation:

$$\left.\frac{\partial\tilde{\mathbf{F}}_{s,j}}{\partial\theta}\right|_{\mathbf{u}_{\mathrm{s}}\text{ fixed}} = \frac{\partial\mathbf{T}_{ms,j}^{\mathrm{T}}}{\partial\theta} \cdot A_j \cdot \sigma_j + \mathbf{T}_{ms,j}^{\mathrm{T}} \cdot \frac{\partial A_j}{\partial\theta} \cdot \sigma_j + \mathbf{T}_{ms,j}^{\mathrm{T}} \cdot A_j \cdot \left.\frac{\partial\sigma_j}{\partial\theta}\right|_{\varepsilon\text{ fixed}} \tag{4}$$

The reason why that is incorrect is that the strain, ε, given to the material, does not only depend on the section deformations, $\mathbf{u}_s$. In fact, the strain is $\varepsilon = \mathbf{T}_{ms}\mathbf{u}_s$, which means that the correct derivative is not Eq. (4) but rather a version with an additional term, here marked with a square bracket parenthesis:

$$\begin{aligned}\left.\frac{\partial\tilde{\mathbf{F}}_{s,j}}{\partial\theta}\right|_{\mathbf{u}_{\mathrm{s}}\text{ fixed}} &= \frac{\partial\mathbf{T}_{ms,j}^{\mathrm{T}}}{\partial\theta} \cdot A_j \cdot \sigma_j + \mathbf{T}_{ms,j}^{\mathrm{T}} \cdot \frac{\partial A_j}{\partial\theta} \cdot \sigma_j + \cdots \\ &\cdots + \left[\mathbf{T}_{ms,j}^{\mathrm{T}} \cdot A_j \cdot \frac{\partial\sigma_j}{\partial\varepsilon_j} \cdot \frac{\partial\varepsilon_j}{\partial\theta}\right] + \mathbf{T}_{ms,j}^{\mathrm{T}} \cdot A_j \cdot \left.\frac{\partial\sigma_j}{\partial\theta}\right|_{\varepsilon\text{ fixed}}\end{aligned} \tag{5}$$

where $\partial\sigma/\partial\varepsilon$ is the material stiffness, such as the modulus of elasticity, E, and, importantly,

$$\frac{\partial\varepsilon}{\partial\theta} = \frac{\partial\mathbf{T}_{\mathrm{ms}}}{\partial\theta} \cdot \mathbf{u}_{\mathrm{s}} \tag{6}$$

Unconditional Derivatives

The second comment is that $\partial\tilde{\mathbf{F}}/\partial\theta\big|_{\mathbf{u}\text{ fixed}}$ actually means $\partial\tilde{\mathbf{F}}/\partial\theta\big|_{\mathbf{u}_{n+1}\text{ fixed}}$. In words, it is the displacements at the current increment that is fixed, not previous increments of the nonlinear analysis. To repeat that point, the direct differentiation method seeks $\partial\tilde{\mathbf{F}}/\partial\theta\big|_{\mathbf{u}_{n+1}\text{ fixed}}$ without any assumption that $\mathbf{u}_n$ and is fixed. For inelastic materials, which employ history variables to calculate the stress for a given trial strain, this necessitates the calculation and storage of unconditional derivatives $\partial\tilde{\mathbf{F}}/\partial\theta$. Using the jargon "Phase 1" and "Phase 2" from Zhang and Der Kiureghian (1993), this means that the materials must first return $\partial\sigma/\partial\theta|_{\varepsilon\text{ fixed}}$ in Phase 1, in order to calculate $\frac{\partial\tilde{\mathbf{F}}}{\partial\theta}\Big|_{\mathbf{u}\text{ fixed}}$, and then later calculate and store $\partial\sigma/\partial\theta$, as well as potentially additional derivatives of history variables, once $\partial\varepsilon/\partial\theta$ is available from the calculation of $\partial\mathbf{u}/\partial\theta$.